Compressive direct imaging of a billion-dimensional optical phase space

Samuel H. Knarr* and Daniel J. Lum
Department of Physics and Astronomy, University of Rochester, Rochester, New York 14627, USA

James Schneeloch
Air Force Research Laboratory, Information Directorate, Rome, New York 13441, USA

John C. Howell
Department of Physics and Astronomy, University of Rochester, Rochester, New York 14627, USA
and Racah Institute of Physics, The Hebrew University of Jerusalem, Jerusalem 91904, Givat Ram, Israel

(Received 12 June 2018; published 28 August 2018)

Optical phase spaces represent fields of any spatial coherence and are typically measured through phase-retrieval methods involving a computational inversion, optical interference, or a resolution-limiting lenslet array. Recently, a weak-values technique demonstrated that a beam's Dirac phase space is proportional to the measurable complex weak value, regardless of coherence. These direct measurements require raster scanning through all position-polarization couplings, limiting their dimensionality to less than 100 000 [C. Bamber and J. S. Lundeen, Phys. Rev. Lett. **112**, 070405 (2014)]. We circumvent these limitations using compressive sensing, a numerical protocol that allows us to undersample, yet efficiently measure, high-dimensional phase spaces. We also propose an improved technique that allows us to directly measure phase spaces with high spatial resolution with scalable frequency resolution. With this method, we are able to easily and rapidly measure a 1.07-billion-dimensional phase space. The distributions are numerically propagated to an object in the beam path, with excellent agreement for coherent and partially coherent sources. This protocol has broad implications in quantum research, signal processing, and imaging, including the recovery of Fourier amplitudes in any dimension with linear algorithmic solutions and ultra-high-dimensional phase-space imaging.

DOI: 10.1103/PhysRevA.98.023854

I. INTRODUCTION

Phase-space representations of light are typically functions of conjugate variables allowing the description of full optical fields of any coherence [1–3]. This information has applications in lensless imaging [3] and beam shaping [4,5] as well as imaging in scattering media [6]. While measuring a spatially coherent beam's amplitude and phase through well-established techniques such as optical interference [7] or phase retrieval [8] is relatively straightforward, tomographical measurements of partially coherent beams at high resolution are a laborious challenge, and many require a computational inversion [3,9] or imprecisely scanning physical components to recover the phase-space distribution. Fortunately, tomography is a standard tool in quantum research used for estimating quantum states [10,11], and the language of quantum mechanics allows us to develop new tools even for classical fields [12].

Recently, a new tomography method was introduced using quantum weak-value techniques [13–15] to directly measure physical states without optical interference or numerical inversion. Unlike typical quantum tomographical methods that estimate states in terms of the density matrix ρ [11] or Wigner function [3,10], the simplest form of the weak-value tomography measures the Dirac phase space, also known as the Kirkwood-Rihaczek distribution [16–18]. The Dirac phase space is a non-Hermitian complex quasiprobability distribution related to the Fourier transform of the density matrix [19,20]. When describing quasimonochromatic and stationary light, optical phase spaces are functions of four transverse variables: two spatial coordinates and two spatial frequencies. Explicitly, if we have a system with density matrix ρ with transverse positions $\mathbf{x} = (x, y)$ and spatial frequencies $\mathbf{k} = (k_x, k_y)$, the antistandard ordered Dirac representation [21] is

$$S(\mathbf{x}, \mathbf{k}) = \mathrm{Tr}[|\mathbf{k}\rangle\langle\mathbf{k}|\mathbf{x}\rangle\langle\mathbf{x}|\rho] = \langle\mathbf{k}|\mathbf{x}\rangle\langle\mathbf{x}|\rho|\mathbf{k}\rangle, \quad (1)$$

while the standard ordered distribution is

$$S(\mathbf{k}, \mathbf{x}) = \mathrm{Tr}[|\mathbf{x}\rangle\langle\mathbf{x}|\mathbf{k}\rangle\langle\mathbf{k}|\rho] = \langle\mathbf{x}|\mathbf{k}\rangle\langle\mathbf{k}|\rho|\mathbf{x}\rangle, \quad (2)$$

where $\mathrm{Tr}[*]$ is the trace. While these distributions are simply complex conjugates, their measurement sequences are different. As with the Wigner function, the marginals, taken by summing over position or frequency variables, give the positive intensity and spectrum of the field: $\langle\mathbf{x}|\rho|\mathbf{x}\rangle$ and $\langle\mathbf{k}|\rho|\mathbf{k}\rangle$.

The idea behind these weak-values tomographies is to weakly couple the preselected state of interest (like the transverse state) to another independent degree of freedom, a meter state (such as polarization), and then filter the result by specific measurement outcomes (postselection) in a basis conjugate to the degree of freedom of the preselected state (the Fourier plane of the preselected plane). For example, in the measurement of (1), analysis of the meter state gives the complex weak

*sknarr@ur.rochester.edu

value [15]

$$A_w = \frac{\langle \mathbf{k}|\mathbf{x}\rangle\langle \mathbf{x}|\rho|\mathbf{k}\rangle}{\langle \mathbf{k}|\rho|\mathbf{k}\rangle} \sim S(\mathbf{x},\mathbf{k}) \tag{3}$$

with measurable real and imaginary parts. This method allows for direct measurement of the phase space, removing the need for a computational inversion. Experimental demonstrations of direct measurements of (1) were first performed measuring transverse pure states [14,22] and later mixed (incoherent) states [21]. Weak-value tomographies have also been performed on polarization [23,24] and orbital angular momentum states [25]. Weak-value direct measurements have also been used to find the quantum process matrix of an arbitrary operation [26].

Direct tomography, while reducing computational complexity, does not necessarily reduce the number of required measurements. Previous demonstrations required scanning through all possible measurements, limiting tests to low-dimensional systems. Here we circumvent previous limitations by incorporating compressive sensing (CS) [27–29], which allows us to reduce the resource requirements, especially for high-dimensional states. CS is a numerical method that reconstructs undersampled signals after sampling in a compressive way. CS relies on the assumption of sparsity or approximate sparsity, i.e., the signal of interest has few (or few significant) nonzero components in a predefined basis. This assumption works quite well for most signals, as signals of interest usually have structure in some basis that set them apart from random noise. CS has also shown promise in quantum systems [30–32] and in the coupling interaction in weak-value wavefront sensing [33,34].

In this work, we improve upon previous tests by compressively measuring high-dimensional classical antistandard Dirac spaces. In CS terms, we are using many low-dimensional single-pixel cameras to recover the higher dimensional phase space. Our method goes significantly beyond past incorporations of CS into weak-value measurements [33,34], in that it works in the more general case of partially coherent light. We then modify this method to measure the standard Dirac phase space, allowing for faster acquisition of even larger phase spaces. With this technique, we easily and efficiently measure phase spaces of more than one billion dimensions. Additionally, we expand on previous tests by using a strong polarization-position coupling, which mitigates state-estimation errors [35].

II. THEORY

A. Imaging the antistandard Dirac space compressively

To introduce CS into the measurement of (1), we use the experimental setup shown in Fig. 1. The light passes through an $\hbar$ cutout and propagates freely to a polarizer to prepare the meter state. Since we are interested in the transverse distribution, we couple position to polarization by using a spatial light modulator (SLM) to rotate the polarization by an angle θ at certain pixels. We apply M masks of $N \times N$ random binary patterns $f_i(\mathbf{x})$, for $i = 1, 2, \ldots, M$, to the SLM, which couple random positions with polarization. As the measurement is compressive, $M \ll N^2$. The SLM's operation

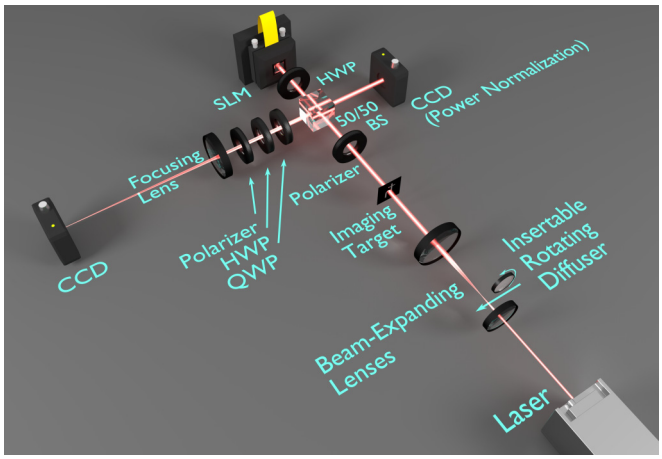

FIG. 1. Experimental setup: Light passes through an $\hbar$ paper cutout before propagating 40 cm to a horizontal polarizer and SLM. To make the source partially coherent, we insert a rotating glass diffuser in the focal plane of the beam expander. The SLM rotates the polarization at randomly chosen pixels to couple position and polarization. After passing through polarization projection optics, a lens Fourier transforms the light reflected off the SLM onto a postselection camera. A separate camera acts as a normalizing bucket detector on the reflected port of the beam splitter in front of the SLM.

is

$$\hat{U}(f_i(\mathbf{x}),\theta) = \int d\mathbf{x}\,\{\mathbb{1}_\sigma[1 - f_i(\mathbf{x})] + f_i(\mathbf{x})e^{-i\sigma\cdot\mathbf{n}\theta}\}|\mathbf{x}\rangle\langle\mathbf{x}|, \tag{4}$$

where σ is the usual Pauli operator, $\mathbb{1}_\sigma$ is the Pauli identity, and $\mathbf{n}$ is the axis of rotation. The first term ensures pixels where $f_i(\mathbf{x}) = 0$ stay horizontally polarized, while the second term rotates the polarization of pixels where $f_i(\mathbf{x}) = 1$. For the strongest possible coupling, we let $\theta \to \pi/2$, so these pixels are rotated to vertical polarization. This operation is shown in Fig. 2.

For simplicity, let the input state be a pure state $|\Psi\rangle = |\psi(\mathbf{x})\rangle|H\rangle$. Applying the operation in Eq. (4), the input state becomes

$$U(f_i(\mathbf{x}),\theta)|\Psi\rangle = \int d\mathbf{x}\psi(\mathbf{x})\{[1 - f_i(\mathbf{x})]|H\rangle + f_i(\mathbf{x})|V\rangle\}|\mathbf{x}\rangle. \tag{5}$$

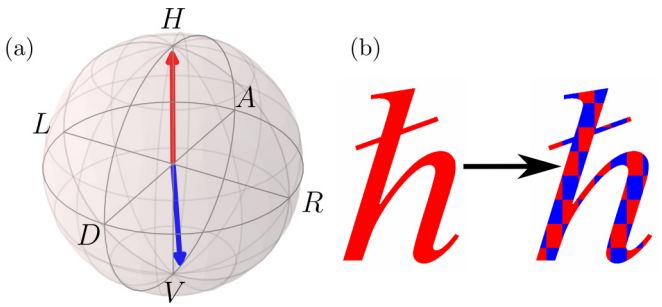

FIG. 2. Polarization projections: (a) Poincaré sphere showing the nearly orthogonal polarizations that each pixel occupies as the light leaves our SLM. (b) An $\hbar$ pixelated into the two outgoing polarizations.

Next, we postselect on different frequency components. This postselection creates an interference between the position and frequency projections of ρ that allows for transverse phase-information retrieval. To postselect, a lens Fourier-transforms the light reflected off of the SLM onto a camera, which takes $N \times N$ pictures for each coupling mask on the SLM. The unnormalized state at each camera pixel is

$$|\psi_{ps}\rangle \sim \tilde{\psi}(\mathbf{k})|H\rangle - \int d\mathbf{x} f_i(\mathbf{x})\psi(\mathbf{x})e^{-i\mathbf{k}\cdot\mathbf{x}}(|H\rangle - |V\rangle), \tag{6}$$

where $\tilde{\psi}(\mathbf{k})$ is the Fourier transform of $\psi(\mathbf{x})$. Note that the postselection is a pixelwise operation, meaning that there is a strong frequency projection measurement at each pixel. Since the transverse state is coupled to polarization, a polarization analysis (we record an image for each mask of each polarization projection) determines the real and imaginary parts of the transverse Dirac distribution.

For a pure state, the Dirac distribution (1) can be written as $S(\mathbf{x}, \mathbf{k}) \sim \psi(\mathbf{x})\tilde{\psi}^*(\mathbf{k})e^{-i\mathbf{k}\cdot\mathbf{x}}$ where ψ^* is the complex conjugate of ψ. By performing the following polarization projections [35], we can measure the partially compressed real and imaginary parts of the Dirac distribution at each pixel (k_x, k_y):

$$y_{\text{real},i,\mathbf{k}} = \int d\mathbf{x} f_i(\mathbf{x})\text{Re}\{\text{S}(\mathbf{x}, \mathbf{k})\} \sim \langle \sigma_x + 2|V\rangle\langle V|\rangle_{\psi_{ps}},$$
$$y_{\text{imag},i,\mathbf{k}} = -\int d\mathbf{x} f_i(\mathbf{x})\text{Im}\{\text{S}(\mathbf{x}, \mathbf{k})\} \sim \langle \sigma_y\rangle_{\psi_{ps}}. \tag{7}$$

To more efficiently measure the non-Hermitian operator $\sigma_x + i\sigma_y$, we reduced our number of measurements by decomposing this into a complex sum of Hermitian operators as in [36]

$$\sigma_x + i\sigma_y = \frac{4}{3}\sum_{q=0}^{2} e^{2\pi i q/3}|s_q\rangle\langle s_q|, \tag{8}$$

where

$$|s_q\rangle = \frac{1}{\sqrt{2}}(|H\rangle + e^{-4\pi i q/3}|V\rangle). \tag{9}$$

If we think of our four-dimensional Dirac phase space as a two-dimensional (2D) space with indexed positions (x, y) as rows and spatial frequencies (k_x, k_y) as columns, the combined images of the polarization projections for a single coupling mask correspond to a single row. That is, we measure an $M \times N^2$ compressed phase space, where each of the M rows corresponds to a coupling mask. Essentially in this 2D reshaping picture, the phase space is compressed in only one dimension. Each column is then separately reconstructed from the measurement matrix. Thus, reconstruction of the full phase space comes from solving many smaller CS problems.

B. Standard Dirac phase space

To modify this experiment to measure (2), we insert a lens after the object to focus the light onto the SLM. This effectively reverses the domains of the previous test; the spatial frequencies are coupled to polarization, and the postselection camera captures the transverse positions. All of the previous analysis applies with the appropriate changes. For example, we now write the coupling interaction of the SLM as

$$\hat{U}(f_i(\mathbf{k}), \theta) = \int d\mathbf{k} \left\{\mathbb{1}_\sigma[1 - f_i(\mathbf{k})] + f_i(\mathbf{k})e^{-i\sigma\cdot\mathbf{n}\theta}\right\}|\mathbf{k}\rangle\langle\mathbf{k}|. \tag{10}$$

The compressed real and imaginary components at each camera pixel (x, y) then are

$$y_{\text{real},i,\mathbf{x}} = \int d\mathbf{k} f_i(\mathbf{k})\text{Re}\{\text{S}(\mathbf{k}, \mathbf{x})\} \sim \langle \sigma_x + 2|V\rangle\langle V|\rangle_{\psi_{ps}},$$
$$y_{\text{imag},i,\mathbf{x}} = -\int d\mathbf{k} f_i(\mathbf{k})\text{Im}\{\text{S}(\mathbf{k}, \mathbf{x})\} \sim \langle \sigma_y\rangle_{\psi_{ps}}. \tag{11}$$

While this is a simple experimental change, it offers several practical advantages. First, we are now able to use the high-resolution postselection camera to gather more spatial dimensions. Since most practical applications do not require such high-resolution spatial frequency information, we can perform faster lower resolution CS scans on the SLM (see below for data collection benefits). These lower dimensional CS reconstructions also lower the computational burden, while still acquiring a higher dimensional phase space faster than the previous method.

III. CS RECONSTRUCTIONS

We reconstructed the real and imaginary parts of the Dirac phase space separately, one component at a time. For our CS solver, we use a total-variation-minimization solver TVAL3 [37], which searches for solutions to the problem

$$\arg\min_{\mathbf{x}} \left[||\mathbf{D}\mathbf{x}||_1 + \frac{\mu}{2}||\mathbf{A}\mathbf{x} - \mathbf{y}||_2^2\right], \tag{12}$$

where $\mathbf{D}$ is the discrete gradient across $\mathbf{x}$ and $||*||_p$ is the L^p norm. Variable $\mathbf{A}$ is a matrix containing our projectors, and variable $\mathbf{y}$ is a vector containing our measurement results. By using this solver, we assume that the total variation of each distribution is sparse across the SLM. We make this assumption because TVAL3 has been shown to work well in image processing; images are often well defined by their edges, and an image's gradient will emphasize edges. It is not unreasonable to guess that it should work in a similar way here. Beyond this, we do not make any assumptions about the sparsity of the Dirac distribution and performed all reconstructions in the pixel basis. Unfortunately, this algorithm did not incorporate any of the physical constraints [20] on the Dirac distributions, which left unphysical artifacts in the reconstructions.

Next, we performed a Bayesian shrinkage denoising algorithm with soft thresholding [38] followed by a low (1%) hard thresholding to remove low-level high-frequency noise from the Dirac distribution. With TV minimization identifying the significant signal components and finding their quantity to be less than the number of measurements, we then further correct the phase-space elements using a least-squares fitting algorithm on each reconstruction result that assumed Gaussian noise in the signal [39].

IV. EXPERIMENT

In our first experiment measuring the antistandard Dirac space, our source is a helium-neon (HeNe) laser with and without a rotating ground glass diffuser for coherent and partially coherent illumination. We use $N = 128$, so that the measured phase-space dimensionality is 268 million (i.e., 128^4), and we use randomized Sylvester-Hadamard patterns, which are composed of 1's and -1's. To perform this operation, we split each pattern into a projection with 1's and 0's, and subtract an inverse pattern with the 0's and 1's switched. The coupling interaction uses a Meadowlark Optics XY spatial light modulator (512×512 pixels with a pixel pitch of 15 μm, binned into 128×128 pixels). We calibrate our system through a polarization tomography [11] so that the unrotated state is $|H\rangle$, and the rotated state is as close to orthogonal ($\theta \approx 0.98\pi$) as possible, such that we can reasonably approximate it as $|V\rangle$.

We Fourier transform the light on the face of the SLM using a 250 mm lens onto a Thorlabs DCC1545M camera (1280×1024 pixels with 5.2 μm $\times$ 5.2 μm pixels; we use the center 512×512 pixel section binned into 128×128 superpixels). Unfortunately, the light quickly saturates the camera since it is in the focal plane. This means that we have to attenuate the light and average several images per pattern in order to resolve the high-frequency components. For the coherent illumination, we average 16 images taken with 6 ms integration times, while for the partially coherent light, we average 64 of these images. The pixels report 8-bit intensity values; to remove background noise, we threshold away any pixel value less than 1 bit. All images are normalized for power fluctuations by using the bucket detector shown in Fig. 1.

The marginals of the measured distributions are displayed in Fig. 3; these should be the probability distributions of the light on the SLM and camera. To remove any remaining unphysical negative and imaginary values, we show the real part with negative values thresholded to zero. As expected after the long free space propagation, we see blurry images of the $\hbar$ and their diffraction patterns. In the coherent source marginals, we can clearly see fringes on the $\hbar$ and a tight diffraction profile. For the partially coherent light, the $\hbar$ is simply blurred, and the diffraction pattern is quite broad. We perform this measurement using a 20% sampling rate. That is, the measurement uses approximately 3300 out of 128^2 possible coupling projections, recording an image for each pattern and polarization projection.

To test the reconstruction's accuracy, we numerically propagate the reconstructions 40 cm back to the $\hbar$ cutout to recover the light at the object. This is done by solving the four-dimensional Bayesian propagation integral [20]

$$S(\mathbf{x}', \mathbf{k}') = \int d\mathbf{x}\, d\mathbf{k}\, K(\mathbf{x}, \mathbf{k}; \mathbf{x}', \mathbf{k}') S(\mathbf{x}; \mathbf{k}), \qquad (13)$$

where $K(\mathbf{x}, \mathbf{k}; \mathbf{x}', \mathbf{k}')$ is the Dirac phase-space propagator. Since the propagation is through free space, the spatial frequency integrals vanish leaving an integral that can be easily evaluated with standard computational Fourier methods, independent of the spatial coherence of the source. After propagating the distributions back to the object plane, we find sharp object images shown in Fig. 4. The position marginals of the propagated distributions are in excellent agreement with

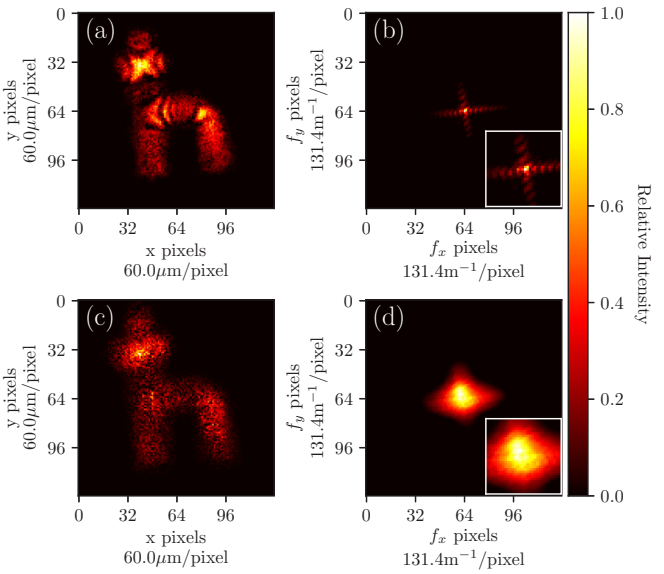

FIG. 3. Marginals of the measured antistandard Dirac distributions: The top row shows the results from the coherent source; the bottom row contains the marginals from the partially coherent source. All probability distributions have been individually normalized to the same scale for visual clarity. (a), (c) The position marginals corresponding to the intensities on the SLM. (b), (d) The frequency marginals, here in units $\mathbf{k}/2\pi$, showing the diffraction patterns on the camera. The insets show closer views of the center 64×64 pixels of the diffraction patterns.

an image taken of the cutout. The partially coherent results contain higher levels of background noise, but this is most likely due to the weaker signal after the diffuser. Again we remove unphysical values by showing the real positive part of the marginals.

To measure the standard Dirac phase space, we use the same configuration and data-processing steps as above except with a 250 mm focal length lens inserted before the SLM such that the SLM is in the back focal plane. We switch our simple cutout for a DLP Lightcrafter 3000 digital micromirror device (DMD) displaying a crossed three-slit pattern. In this configuration, an image of the front focal plane of the new lens now appears on the camera as it is a $4f$ imaging system. We displace the DMD 10 cm from this front focal plane, severely blurring the image of the object on the camera. We broadly illuminate this object with partially coherent radiation from a red LED, frequency filtered with a 633 nm line filter to work with the SLM. An iris was placed just before the SLM to block the light from reflecting off the metal edges, which would only add noise to the signal.

We place 64×64 coupling patterns on the SLM and take 512×512 images on the postselection camera. Thus we acquire a 1.07-billion-dimensional phase space (512^2 spatial dimensions and 64^2 spatial frequencies). Our measured and propagated results are shown in Fig. 5 for a 20% sample rate. The measured position marginal [Fig. 5(b)] corresponding to the light on the camera barely has any recognizable structure to it, and yet clearly after propagation [Fig. 5(c)] we can see a sharp three-slit object. This altered setup has several practical benefits that shift the required resources from experiment to

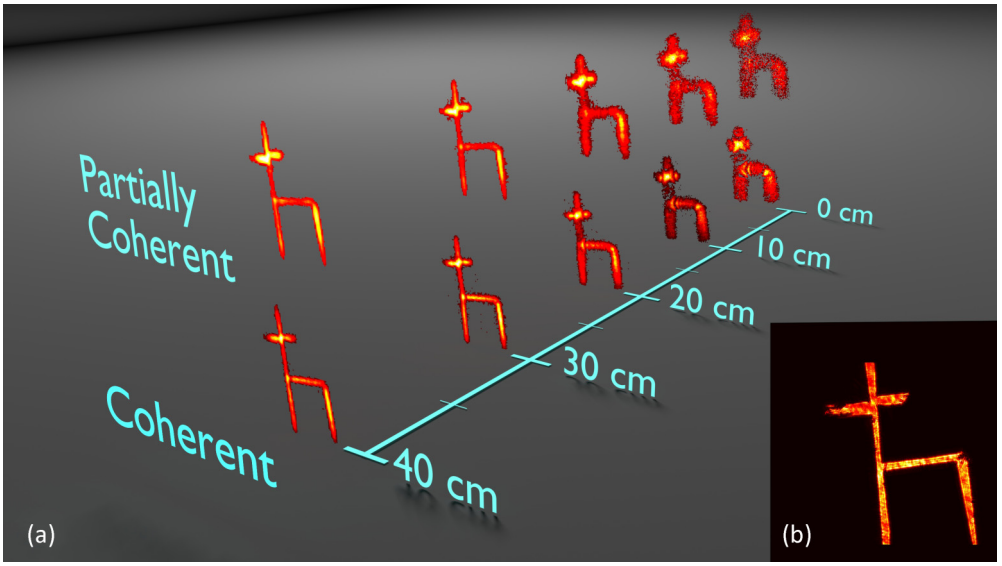

FIG. 4. Comparison of numerical propagation with actual image: Plots have been normalized to be shown on the same scale. (a) The position marginal of the propagated Dirac distributions showing the intensities as functions of distance from the SLM (0 cm) to the object (40 cm). The upper row shows the partially coherent light, while the lower row shows the coherent light. (b) Actual image of the object. The propagation data are in good agreement with the image for both sources.

computation. First, since the camera is now in an image plane of the system, we avoid the saturation effects affecting the previous scheme. Thus we can take single images, removing the need for averaging. For this test, we use a 25 ms integration time (due to the dimness of the light on the camera).

Since we are reconstructing the Dirac phase space's real and imaginary part at each camera pixel separately, theoretically we require $2N^2$ reconstructions. However, when measuring the antistandard ordered Dirac matrix, the measurement matrices $\mathbf{y}$ are quite sparse as many of the spatial frequencies are zeros, so no reconstruction is needed for these components since there is no signal. This is confirmed in Figs. 3(b) and 3(d). In total, we have to perform only 4544 reconstructions for the coherent illumination and 2871 reconstructions for the partially coherent light of 128×128 pixel images. We would normally expect that the partially coherent light requires more reconstructions than the coherent light since its measurement vector should have more frequency components. However, it is likely that the weaker high-frequency components are too dim for the camera to see above the background noise. This is not quite the case in measuring the standard ordered Dirac distribution, which requires more than 200 K reconstructions for each the real and imaginary due to the broad image distribution on the camera. However, these are smaller 64×64 reconstructions, which can be done quite quickly on modern computers. Furthermore, since each reconstruction is independent, parallelized operations are used to significantly reduce processing time.

V. DISCUSSION

Our new method measuring the standard Dirac distribution experiment significantly outperforms the compressed antistandard test. For the lower-dimensional phase-space antistandard test, the total time spent collecting light is 1.5 h for the coherent illumination and 5.5 h of signal collection for partially coherent, whereas in the high-resolution standard Dirac test, it is only 2.8 min. The reduction in collection times is due to the standard Dirac experiment avoiding the need for

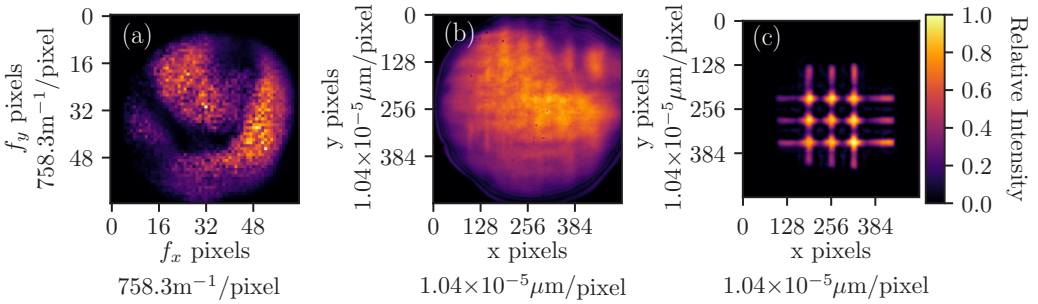

FIG. 5. Measurement of the standard ordered Dirac distribution: Plots have been normalized to be shown on the same scale. (a) The frequency marginal of the Dirac distributions showing the intensities across the SLM. The circular outline is due to an iris blocking the light from the highly reflective edges of the SLM. (b) Spatial marginal corresponding to the extremely out-of-focus image on the camera. (c) The spatial marginal after propagation, clearly showing the well-defined object.

image averaging (see Methods). Instead, it shifts the important high-resolution spatial information from the slower SLM CS acquisition to the faster camera. The actual experiment time is much longer than the integration times, but fortunately this is almost entirely due to overhead from equipment limitations. For example, most of this extra time came from the cameras not having internal storage and requiring each image acquisition be transferred to a computer. Additionally, our SLM required 100 ms to stabilize for each coupling pattern. With additional resources and better equipment this overhead time could be virtually eliminated and could even reduce the time spent collecting light. By performing the polarization projections and 1's and −1's CS projections in parallel, we would reduce the light collection time by a factor of 8, meaning that this measurement could be done in seconds.

Also, by using a postselection camera combined with compressive coupling, our required measurements scale only as the resolution of our SLM N^2, while measuring an N^4 dimensional phase space. Thus our technique allows rapid acquisition of these high-dimensional phase spaces. It also handles the high-dimensional phase-space data in a computationally efficient way. Reconstructions and least-square fittings are performed component-wise independently, limiting the number of operations necessary on the full distribution.

Our demonstration here uses relatively simple scenes, and our reconstructions assume only that total variation of the beam across the SLM pixels would be sparse. However, total variation has been shown to work well for natural images, and in practice one would usually know something about the signal of interest, allowing them to choose a basis for reconstruction where the signal should have a sparse representation. For example, natural images are known to also have sparse representations in the discrete cosine and wavelet bases. Also, as previously stated, the TVAL3 solver does not incorporate any physical constraints [20] on the reconstructions, and so experimental noise left unphysical artifacts in the phase space. The denoising and least squares steps fix some of the values, but the result is still slightly unphysical. This results in marginals having complex and negative values when they should be entirely positive since they are probabilities. From a practical point of view, this does not strongly affect the propagation and refocusing of the light, and by viewing the real positive part of the marginals, we have good agreement in the data shown in Figs. 4 and 5.

VI. CONCLUSION

We have shown that we can efficiently and rapidly acquire high-dimensional optical phase spaces for light of any spatial coherence. With this information, we can numerically propagate the light to any plane for lensless imaging. This goes significantly further than previous demonstrations that worked in only one dimension and required scanning at every coupling pixel [21]. Unlike other phase-space techniques, we did not need to use lenslet arrays or moving components, and we were able to directly measure the phase-space elements. Another benefit to this method is that the measured Dirac distribution scales as the product of the number of pixels in the coupling interaction with the number of pixels in the postselection, meaning that we can very easily measure extremely high-dimensional phase spaces. In the future, customized algorithms incorporating physical constraints could give better reconstructions and further reduce the number of measurements and the numerical resources required.

ACKNOWLEDGMENTS

S.H.K. thanks G. A. Howland for insightful discussions. S.H.K, D.J.L, and J.C.H acknowledge support from the Air Force Office of Scientific Research, Grant No. FA9550-16-1-0359, and from Northrop Grumman, Grant No. 058264-002. J.S. acknowledges support from the National Research Council Research Associate Programs and funding from the OSD ARAP QSEP program. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of AFRL.

S.H.K conceived the idea and performed the experiment. D.J.L. provided the algorithms for data analysis. J.S. provided theoretical support. J.C.H. supervised the project. S.H.K. prepared the manuscript with contributions from all other authors.

APPENDIX: PROPAGATION

As noted in Refs. [15,20], the Dirac distribution can be propagated in a Bayesian manner. Therefore, using the definition of the Dirac distribution [Eq. (1)], for $S(\mathbf{x}, \mathbf{k})$ we can find the Dirac representation in the plane (x', y'), $S(\mathbf{x}', \mathbf{k}')$, through

$$\begin{aligned} S(\mathbf{x}', \mathbf{k}') &= \int d\mathbf{x}\, d\mathbf{k}\, \langle \mathbf{k}'|\mathbf{x}'\rangle \langle \mathbf{x}'|\mathbf{x}\rangle \langle \mathbf{x}|\rho|\mathbf{k}\rangle \langle \mathbf{k}|\mathbf{k}'\rangle \\ &= \int d\mathbf{x}\, d\mathbf{k}\, \frac{\langle \mathbf{k}'|\mathbf{x}'\rangle \langle \mathbf{x}'|\mathbf{x}\rangle \langle \mathbf{k}|\mathbf{k}'\rangle}{\langle \mathbf{k}|\mathbf{x}\rangle} S(\mathbf{x}; \mathbf{k}), \end{aligned} \quad \text{(A1)}$$

where the propagator $K(\mathbf{x}, \mathbf{k}; \mathbf{x}'\mathbf{k}')$ of Eq. (13) is immediately identified. For free-space propagation over a distance z, the terms in K simplify to

$$\begin{aligned} \frac{\langle \mathbf{k}'|\mathbf{x}'\rangle}{\langle \mathbf{k}|\mathbf{x}\rangle} &= e^{i(\mathbf{k}\cdot\mathbf{x} - \mathbf{k}'\cdot\mathbf{x}')}, \\ \langle \mathbf{k}|\mathbf{k}'\rangle &= e^{-iz\sqrt{k^2 - k_x^2 - k_y^2}} \delta(\mathbf{k} - \mathbf{k}'), \\ \langle \mathbf{x}'|\mathbf{x}\rangle &= \frac{1}{(2\pi)^2} \int d\mathbf{k}''\, e^{i\mathbf{k}''\cdot(\mathbf{x} - \mathbf{x}')} e^{iz\sqrt{k^2 - k_x''^2 - k_y''^2}}, \end{aligned} \quad \text{(A2)}$$

where the middle term shows that through free-space propagation, momentum is conserved. Putting this all together and rearranging terms results in the propagation equation

$$\begin{aligned} S(\mathbf{x}, \mathbf{k}) = \frac{1}{(2\pi)^2} e^{-iz\sqrt{k^2 - k_x^2 - k_y^2}} e^{-i\mathbf{k}\cdot\mathbf{x}}\, \mathrm{F}^{-1}_{\mathbf{k}''\to\mathbf{x}} \\ \times \{ e^{iz\sqrt{k^2 - k_x''^2 - k_y''^2}}\, \mathrm{F}_{\mathbf{x}'\to\mathbf{k}''} [e^{i\mathbf{k}\cdot\mathbf{x}'} S(\mathbf{x}', \mathbf{k})] \}, \end{aligned} \quad \text{(A3)}$$

where F (F^{-1}) indicates a (inverse) Fourier transform. This equation is very similar to Fourier optics field propagation equations and can be easily evaluated with numerical Fourier methods. However, (A3) works for any field regardless of coherence. Also note that operationally for a discrete phase space, if we reshape it into a 2D distribution $(\mathbf{x}, \mathbf{k})$, we are propagating each spatial frequency column separately.

To propagate (2), we could follow a similar procedure as above, first propagating the distribution from the SLM to the lens, through the lens, and then to the object. However, we

simply make the change $\rho \rightarrow U\rho U^{\dagger}$, such that

$$U = \exp\left(if\frac{k_x^2+k_y^2}{2k}\right)\exp\left(-ik\frac{x^2+y^2}{2f}\right)\exp\left(iz\frac{k_x^2+k_y^2}{2k}\right), \tag{A4}$$

where z is the distance from the object to the lens and f is the lens focal length. We then insert this into (2). Unlike the propagation above, the spatial frequency integrals do not vanish, but standard Fourier propagation methods are still used to find the distribution at the object.

[1] M. A. Alonso, Adv. Opt. Photonics **3**, 272 (2011).
[2] L. Mandel and E. Wolf, *Optical Coherence and Quantum Optics* (Cambridge University Press, Cambridge, 1995).
[3] L. Waller, G. Situ, and J. W. Fleischer, Nat. Photon. **6**, 474 (2012).
[4] O. Katz, E. Small, Y. Bromberg, and Y. Silberberg, Nat. Photon. **5**, 372 (2011).
[5] A. P. Mosk, A. Lagendijk, G. Lerosey, and M. Fink, Nat. Photon. **6**, 283 (2012).
[6] H.-Y. Liu, E. Jonas, L. Tian, J. Zhong, B. Recht, and L. Waller, Opt. Express **23**, 14461 (2015).
[7] I. Yamaguchi and T. Zhang, Opt. Lett. **22**, 1268 (1997).
[8] J. R. Fienup, Appl. Opt. **21**, 2758 (1982).
[9] K. A. Sharma, T. G. Brown, and M. A. Alonso, Opt. Express **24**, 16099 (2016).
[10] A. I. Lvovsky and M. G. Raymer, Rev. Mod. Phys. **81**, 299 (2009).
[11] D. F. V. James, P. G. Kwiat, W. J. Munro, and A. G. White, Phys. Rev. A **64**, 052312 (2001).
[12] B. Stoklasa, L. Motka, J. Rehacek, Z. Hradil, and L. Sánchez-Soto, Nat. Commun. **5**, 3275 (2014).
[13] J. Dressel, M. Malik, F. M. Miatto, A. N. Jordan, and R. W. Boyd, Rev. Mod. Phys. **86**, 307 (2014).
[14] J. S. Lundeen, B. Sutherland, A. Patel, C. Stewart, and C. Bamber, Nature (London) **474**, 188 (2011).
[15] J. S. Lundeen and C. Bamber, Phys. Rev. Lett. **108**, 070402 (2012).
[16] P. A. M. Dirac, Rev. Mod. Phys. **17**, 195 (1945).
[17] L. M. Johansen, Phys. Rev. A **76**, 012119 (2007).
[18] V. Bollen, Y. M. Sua, and K. F. Lee, Phys. Rev. A **81**, 063826 (2010).
[19] S. Chaturvedi, E. Ercolessi, G. Marmo, G. Morandi, N. Mukunda, and R. Simon, J. Phys. A **39**, 1405 (2006).
[20] H. F. Hofmann, New J. Phys. **14**, 043031 (2012).
[21] C. Bamber and J. S. Lundeen, Phys. Rev. Lett. **112**, 070405 (2014).
[22] Z. Shi, M. Mirhosseini, J. Margiewicz, M. Malik, F. Rivera, Z. Zhu, and R. W. Boyd, Optica **2**, 388 (2015).
[23] J. Z. Salvail, M. Agnew, A. S. Johnson, E. Bolduc, J. Leach, and R. W. Boyd, Nat. Photon. **7**, 316 (2013).
[24] G. S. Thekkadath, L. Giner, Y. Chalich, M. J. Horton, J. Banker, and J. S. Lundeen, Phys. Rev. Lett. **117**, 120401 (2016).
[25] M. Malik, M. Mirhosseini, M. P. Lavery, J. Leach, M. J. Padgett, and R. W. Boyd, Nat. Commun. **5**, 3115 (2014).
[26] Y. Kim, Y.-S. Kim, S.-Y. Lee, S.-W. Han, S. Moon, Y.-H. Kim, and Y.-W. Cho, Nat. Commun. **9**, 192 (2018).
[27] R. G. Baraniuk, IEEE Sig. Proc. Mag. **24**, 118 (2007).
[28] E. J. Candès and M. B. Wakin, IEEE Sig. Proc. Mag. **25**, 21 (2008).
[29] Y. Shechtman, Y. C. Eldar, O. Cohen, H. N. Chapman, J. Miao, and M. Segev, IEEE Sig. Proc. Mag. **32**, 87 (2015).
[30] G. A. Howland and J. C. Howell, Phys. Rev. X **3**, 011013 (2013).
[31] F. Tonolini, S. Chan, M. Agnew, A. Lindsay, and J. Leach, Sci. Rep. **4**, 6542 EP (2014).
[32] A. Kalev, R. L. Kosut, and I. H. Deutsch, NPJ Quantum Inf. **1**, 15018 (2015).
[33] G. A. Howland, D. J. Lum, and J. C. Howell, Opt. Express **22**, 18870 (2014).
[34] M. Mirhosseini, O. S. Magaña-Loaiza, S. M. H. Rafsanjani, and R. W. Boyd, Phys. Rev. Lett. **113**, 090402 (2014).
[35] G. Vallone and D. Dequal, Phys. Rev. Lett. **116**, 040502 (2016).
[36] E. Bolduc, G. Gariepy, and J. Leach, Nat. Commun. **7**, 10439 (2016).
[37] C. Li, W. Yin, and Y. Zhang, CAAM Report **20**, 46 (2009).
[38] S. G. Chang, B. Yu, and M. Vetterli, IEEE Trans. Image Proc. **9**, 1532 (2000).
[39] R. Fletcher, *Practical Methods of Optimization* (John Wiley & Sons, Hoboken, NJ, 2013).